

PATENT APPLICATION

Compiler and Method for Optimizing Object Codes for Hierarchical Memories

Inventors:

Keiko MOTOKAWA
Citizenship: Japan

Ichiro KYUSHIMA
Citizenship: Japan

Sinichi ITO
Citizenship: Japan

Assignee:

Hitachi, Ltd.
6, Kanda Surugadai 4-chome
Chiyoda-ku, Tokyo, Japan
Incorporation: Japan

Entity:

Large

- 1 -

COMPILER AND METHOD FOR OPTIMIZING OBJECT
CODES FOR HIERARCHICAL MEMORIES

BACKGROUND OF THE INVENTION

The present invention relates to a compiling
5 method capable of reducing execution time of an object
program in computer utilizing techniques. More
specifically, the present invention is directed to an
optimizing-designation method used when a source
program is compiled with respect to an architecture
10 equipped with a plurality of memory hierarchies.

With improvements in operating speeds of
microprocessors, latencies of main storage accesses are
increased. Most of the current processors is provided
with cache memories having relatively small memory
15 capacities, the access speeds of which are faster than
those of the main storage. Furthermore, in some
processors, cache memories are constituted in such a
hierarchical form as a primary cache and a secondary
cache. Since memories are formed in the hierarchical
20 form, data of such a memory hierarchy whose latency is
small is accessed by a processor as many as possible,
so that a total number of accesses by the processor
with respect to data of a memory hierarchy whose
latency is large can be reduced. In other words, a
25 processor which executes a memory access instruction
can access a primary cache in a short latency when the

data hits the primary cache. When the data misses the primary cache, this processor accesses a second cache. When the data hits the secondary cache, the processor can subsequently access data thereof in a short

5 latency. Only when the data misses all of the cache hierarchies, the processor accesses the main storage.

As a method of capable of hiding (namely, not revealing) a latency during cache miss operation, an instruction scheduling method may be employed which
10 executes a process in such a manner that a distance between a load instruction and an instruction (will be referred to as "use instruction" hereinafter) for using the loaded data is made sufficiently longer. A memory latency is employed as a reference for determining how
15 many cycles the distance between both instructions is separated.

A method for "hiding" a latency of referring to the main storage during cache miss operation is described in, for instance, Todd C. Mowry et al,
20 "Design and Evaluation of Compiler Algorithm for Prefetching", Architectural Support for Programming Languages and Operating Systems, pp. 62 to 73 in 1992 (will be referred to as "Publication 1" hereinafter). Publication 1 discloses a so-called "software
25 prefetching (prefetch optimization)". In this software prefetching method, while a prefetch instruction is prepared for a processor and this prefetch instruction instructs to move data from the main storage to a cache

in a preceding manner, a prefetch instruction is inserted into an object program by a compiler. If the prefetch instruction is utilized, then the latency of referring to the main storage can be "hidden". That is, while data to which a processor refers in a succeeding loop iteration is previously moved from the main storage to the cache, another calculation can be carried out by this processor at the same time if this prefetch instruction is utilized.

10 In this software prefetching method, when a prefetch instruction is produced with respect to data reference by the processor within a loop, first of all, the number of execution cycles "C" required for one iteration of this loop is estimated. Next, calculation is made of such a value $\alpha = \text{CEIL}(L/C)$ which is defined by dividing the number of cycles "L" by "C". This cycle number "L" is required in order that data is moved from the main memory to a cache (memory latency). Symbol "CEIL" is assumed as a symbol which indicates
15 rounding up any numbers smaller than, or equal to a decimal point. Since the data to which the processor refers after " α " times of loop iterations has been previously prefetched, when the processor refers to this data after "L" cycles, the data has already been
20 reached to the cache, so that the processor hits the cache and can execute the program at a high speed. In a case where data is prefetched to a primary cache, if the data has already been stored in the primary cache,

then the prefetching operation is no longer required. Also, in a case where data is present in the secondary cache, the number of cycles which requires to move the data from the secondary cache to the primary cache may
5 be used as the memory latency "L", whereas in a case that data is present in the main storage, the number of cycles which requires to move the data from the main storage to the primary cache may be used the memory latency "L". However, it is normally unclear which
10 memory hierarchy data is present in. As a consequence, assuming now that the data is present in the main storage, process is carried out.

Another memory optimizing method is known which can reduce the number of cache misses by way of a
15 program transformation capable of improving a data locality. As a specific program transformation, there are proposed: a loop tiling method, a loop interchanging method, and a loop unrolling method.

The loop tiling method corresponds to a loop
20 transformation operation having the following purpose. That is, in a case where data to which a processor refers within a multiply nested loop owns a reuse, it is so designed that the processor again refers to data which has once been loaded on a cache before this
25 loaded data is ejected from the cache since the processor refers to another data. The loop tiling method is described in Michael Edward Wolf, "Improving Locality and Parallelism in Nested Loops", Technical

Report: CSL-TR-92-538, in 1992 (will be referred to as "Publication 2" hereinafter).

The loop interchange method and the loop unrolling method, which aim to optimize the memory reference pattern, are described in Kevin Dowd, "High Performance Computing", O'Reilly & Associates, Inc., 11.1 section (will be referred to as "Publication 3" hereinafter).

In order to realize the above-described latency hiding optimization and also the above-explained data localization, such information is required which may depend upon attributes of a target machine, for instance, the number of cycles required to memory reference, and a cache size. Normally, information as to a target machine is held as internal information in a compiler. There is another method for instructing the above-described information by a user. In such a publication, i.e., Hitachi, Ltd. (HI-UX/MPP for SR8000) "Optimizing FORTRAN90 User's Guide", 6.2 section (will be referred to as "Publication 4" hereinafter), the following aspect is described. That is, the user can designate that the number of cycles required to read from the memory is "N" by making such option designation as "-mslatency=N" (symbol "N" being positive integer). In a publication, i.e., IBM, "XL Fortran for AIX User's Guide Version 7 Release 1", Chapter 5 (will be referred to as "Publication 5" hereinafter), the following aspect is described. That

is, the user can designate the cache size, the line size, and the associative number every level of the hierarchical cache by making the "-qcache" option.

With respect to the conventional latency
5 hiding optimization and the conventional data
localization, there are different optimizing methods,
depending upon such a condition that data is located in
which memory hierarchy when an object program is
executed.

10 For example, in the instruction scheduling
method, if the distance between a load instruction and
a use instruction is increased, then the total number
of registers to be used is also increased. Also, in
the prefetch optimizing method, when the timing of the
15 prefetch instruction becomes excessively early, there
are some possibilities that the data is again ejected
from the cache before the use instruction is carried
out. As a result, when the memory latency becomes
excessively large which is assumed in these optimizing
20 methods, sufficiently effects achieved by these
optimizing operations cannot be realized. In other
words, when the optimizing operation using the main
storage latency is applied to the data which hits the
L2 cache (secondary cache), there are certain
25 possibilities that the execution performance thereof is
lowered, as compared with that for a case where the
optimizing operation using the L2 cache latency is
applied to the data. However, since it has not clearly

been defined in the prior art which memory hierarchy the subject data of the load instruction is located in, the following problem occurred. That is, the main storage latency had to be assumed to be used when the
5 optimizing operation was applied.

Also, in the data locality optimizing method, if the loop structure is converted into a complex loop structure, overheads of the loop execution will be increased. As a result, there are possibilities that
10 the execution performance is lowered. In a case where the data to which the processor refers within the loop mainly causes the cache miss, the effect may be achieved by applying the data locality optimizing method since the total number of the cache misses is
15 reduced. However, when the cache hits occur, since there is no effect achieved by reducing the total number of the cache misses, it is better not to apply the data locality optimizing method. Since it could not grasp as to whether or not the cache hit occurs in
20 the prior art, the loop transformation has been applied even when the cache hits occur. As a consequence, there has been a problem that the execution performance may be lowered.

SUMMARY OF THE INVENTION

25 To solve the above-described problems, an object of the present invention is to provide a compiler capable of producing object codes optimized by

considering a hierarchical memory structure on a computer system, and also to provide a code producing method of the compiler.

In accordance with the present invention,
5 with respect to a compiler for producing an object program used to be executed on an architecture equipped with a plurality of memory hierarchies from a source program in conjunction with a computer system, both a step for interpreting either an option or a designation
10 statement, and another step for executing an optimizing process directed to the designated memory hierarchy are provided. The option, or the designation statement designates that when a target object program is executed, the target program mainly refers to data
15 present in which memory hierarchy among the plural memory hierarchies.

As the optimizing process directed to the memory hierarchy, a calculating/executing step, or a determining step is provided with the compiler. The
20 former step calculates memory latency according to a designated memory hierarchy in response to a memory access instruction, and then, performs an optimizing process according to the calculated latency. The latter step determines a loop transformation method of
25 loop interchange, loop unrolling, or loop tiling according to a designated memory hierarchy with respect to the memory access instruction.

Fig. 1 schematically indicates an outline of

the present invention. In Fig. 1, assuming now that such data to which a source program "TEST1.f" refers during the execution thereof are mainly L2 cache data and data to which a source program "TEST2.f" refers
5 during the execution thereof are mainly main storage data, optimizing methods with respect to the two source programs should be made different from each other.

However, in the conventional optimizing methods, the optimizing process assuming the main
10 storage data has been applied to both the programs. A memory hierarchy mainly accessed by data to which a program refers has been determined based upon a data size and a loop size. The data size and loop size cannot be cleared by way of a static analysis by a
15 compiler.

In accordance with the present invention, since there is provided means for designating which memory hierarchy data mainly belong to when an object program is executed, a compiler analyzes designation of
20 a memory hierarchy (101), and then, performs an optimizing operation according to the designated memory hierarchy (103, 104, 105). As a result, an object program to which more advanced optimizing operation has been carried out can be produced. In a memory
25 optimizing operation, code conversion is carried out with employment of such attributes as latency and a size of the designated memory hierarchy. In the example of Fig. 1, the optimizing operation (104)

directed to the L2 cache data is applied to the source program TEST1.f, whereas the optimizing operation (105) directed to the main storage data is applied to the source program TEST2.f.

5 In accordance with the present invention, there is provided the means for designating which memory hierarchy data are mainly accessed to in the source program. As a result, application of an optimizing method having no effect when the data is
10 present on the cache can be prevented; the instruction scheduling optimizing method and the prefetch optimizing method with employment of the memory latency according to the designated memory hierarchy can be applied; or the tiling process method with employment
15 of such parameters as a cache size of a target cache according to the designated memory hierarchy can be applied. Whereby, the object program can be executed at a high speed.

Other objects, features and advantages of the
20 invention will become apparent from the following description of the embodiments of the invention taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram for schematically showing
25 an outline of the present invention.

Fig. 2 is a schematic diagram for representing an arrangement of a computer system in

which a compiler for executing an improved optimizing method is operated.

Fig. 3 is a diagram for indicating an example of designation by a compiler option of a memory hierarchy which is mainly accessed when a program is executed.

Fig. 4 is diagram for representing an example of a source program containing a designation statement of a memory hierarchy which is mainly accessed when a program is executed.

Fig. 5 is a flowchart for explaining a processing sequence of a compiler.

Fig. 6 is a diagram for schematically indicating an example of a loop table.

Fig. 7 is a flowchart for explaining a processing sequence of an instruction scheduling method.

Fig. 8 is a diagram for illustratively showing an example of a DAG.

Fig. 9 is a flowchart for describing a sequential process of setting a latency of a DAG edge.

Fig. 10 is a flowchart for describing a sequential process of a prefetch method.

Fig. 11 is a flowchart for describing a sequential process of a tiling method.

Fig. 12 is a flowchart for explaining a sequential process as to a loop interchange and a loop unrolling.

Fig. 13 is a flowchart for describing a sequential process for setting a data designation field.

DESCRIPTION OF THE EMBODIMENTS

5 Referring now to drawings, various embodiments of the present invention will be described.

Fig. 2 is a schematic structural diagram of a computer system in which a compiler according to the present invention is operated.

10 This computer system is arranged by a CPU (central processing unit) 201, a display device 202, a keyboard 203, a main storage apparatus 204, and an external storage apparatus 205. The keyboard 203 accepts a compiler initiating command issued from a
15 user. A compiler end message and an error message are displayed on the display device 202. Both a source program 206 and an object program 207 are stored in the external storage apparatus 205. A compiler 208, an intermediate code 209, and a loop table 210, which are
20 required in a compiling process stage, are stored in the main storage apparatus 204. A compiling process is carried out by that the CPU 201 executes the compiler program 208. It should be noted that the CPU 201 internally has a level-1 (L1) cache (primary cache)
25 2012 and a level-2 (L2) cache (secondary cache) 2013, and constitutes a memory hierarchy in order that process may be executed at a high speed by a process

unit 2011 which contains a fetch/decode unit, an execution unit, and the like. The level-1 cache 2012 owns an access latency different from an access latency of the level-2 cache 2013, which are accessed by this process unit 2011. It is so assumed that in the below-mentioned explanation, such a process implies process/process sequence of a compiler. In this process, the CPU 201 interprets a codes described in a program used in the compiler. It should also be noted that since the computer system shown in Fig. 2 is equipped with the level-1 cache 2012, the level-2 cache 2013, and the main storage apparatus 204, this computer system may be used as such a computer system which executes an object program produced by the compiler according to the present invention.

Fig. 3 indicates an example of a compiler initiating command inputted from a user, which is accepted by the computer system in order to issue a memory hierarchy instruction when a source program is compiled. Symbol "f90" indicates an example of an initiating command of a Fortran compiler, symbol "test.f" shows an example of a designation of a source program, and symbol "-O" indicates an example of a compiler option. In this embodiment, a memory hierarchy designation option is expressed by "-data", and designates that data contained in a program is mainly located in which memory hierarchy. As the memory hierarchies, "L1" indicative of an L1 cache;

"L2" representative of an L2 cache; and "MEM" indicative of a main storage can be designated. Symbol "-data=L2" (301) of Fig. 3 designates that data are mainly present in the L2 cache.

5 In the example of Fig. 3, the memory hierarchy designation is made in the compiler initiating command. The memory hierarchy designation may also be alternatively described as a designation statement contained in a source program. Fig. 4
10 indicates an example of a source program to which a memory hierarchy designation is added. A designation statement 401 designates that data within a loop of "DO 10" mainly hits the L1 cache. Another designation statement 402 designates that data within a loop of "DO
15 20" mainly accesses the main storage. Another designation statement 403 designates that a loop of "DO 30" mainly hits the L2 cache (namely miss L1 cache).

Fig. 5 represent a process sequence of the compiler 208 which is operated in the computer system
20 of Fig. 2. The process of the compiler is carried out in this order of a parsing process 501, a memory optimizing process 502, a register allocating process 503, and a code generating process 504.

In the parsing process 501, while the source
25 program 206 is inputted, both a parsing operation and a loop analyzing operation are carried out so as to output both intermediate codes 209 and a loop table 210. In this parsing step 501, both a memory hierarchy

designation statement and a memory hierarchy designation option are analyzed which constitute the feature of the present invention, and then, the analyzed results are registered in the loop table 210.

5 This loop table 210 will be explained later. Since the memory optimizing process 502 corresponds to a step which constitutes the feature of the present invention, this memory optimizing process 502 will also be described later more in detail. In the register
10 allocating process 503, registers are allocated to the respective nodes of the intermediates codes 209. In the code generating process 504, the intermediate codes 209 are converted into the object program 207 and this object program 207 is outputted.

15 Fig. 6 indicates an example of a content of the loop table 210. In this loop table 210, only a data designation field 602 which constitutes the feature of the present invention is indicated in addition to a loop number 601 used to identify loops of
20 the source program 206. The loop table 210 is produced in the parsing process 501. This table producing process is indicated in a flowchart of Fig. 13, which corresponds to a data designation field setting process.

25 In the sequential process of Fig. 13, loops of the source program 206 are traced in a sequential manner to be processed. In a step 1301, judgement is made as to whether or not a loop which has not yet been

processed is present. When there is a loop which has not yet been processed, an unprocessed loop is derived in a step 1303. In a step 1303, a check is made as to whether or not a "data" designation statement is present in the source program 206 to be compiled. This "data" designation statement corresponds to such a memory hierarchy designation to the loop as shown in Fig. 4. When the "data" designation statement is present in the source program 206, an operand of this "data" designation statement is set in the loop table 210 in a step 1307. In the example of Fig. 4, with respect to the loop of "DO 10", "L1" is set to a column of the "data" designation statement. When the "data" designation statement is not present in the source program 206, the sequential process is advanced to a step 1304. In this step 1304, a check is made as to whether or not a "-data" option corresponding to the memory hierarchy designation option is designated in the compiler initiating command. In a case where the "-data" option is designated in the compiler initiating command, the process is advanced to a step 1305 in which a value of the "-data" option is set. When the "-data" option is not designated in the compiler initiating command, the process is advanced to a step 1306. In this step 1306, "no designation" is set.

Several examples as to the memory optimizing process 502 are described as follows:

As a first example of the memory optimizing

process, an instruction scheduling method is carried out in the memory optimizing process 502. Fig. 7 shows a sequential process of this instruction scheduling method.

5 In a step 701 of the flowchart as to the instruction scheduling method of Fig. 7, a DAG (directed acyclic graph) is formed. Fig. 8 indicates an example of such a DAG corresponding to $A(I)*B(I)+C(I)$. The respective nodes of this DAG
10 correspond to instructions available on intermediate codes. Edges among the nodes indicate restrictions of execution orders. For instance, in a "mul" calculation of a node 803, both a result of "load" of a node 801 and a result of "load" of a node 802 are used as
15 operands, so that both the node 801 and the node 802 must be executed prior to the node 803. This relationship is represented by an edge from the node 801 to the node 803, and by another edge from the node 802 to the node 803.

20 In a step 702, latencies are set to the edges of the DAG formed in the step 701. A detailed process sequence of this step 702 is indicated in Fig. 9.

 In this sequential process, while the edges
25 on the DAG are traced, these edges are sequentially processed. In a step 901 of the flowchart shown in Fig. 9, a judgement is made as to whether or not an edge which has not yet been processed is present. When

there is no unprocessed edge, this process is ended.
To the contrary, when there is such an unprocessed
edge, the process is advanced to a step 902. In this
step 902, the unprocessed edge is derived to be set to
5 an edge which should be processed. In a step 903, a
check is made as to whether or not a starting point
node of the edge corresponds to a load instruction.
When the starting point node is not equal to the load
instruction, the process is advanced to a step 910. In
10 this step 910, a value of a latency corresponding to
the operation of the starting point node is set to the
edge. Then, the process of this edge is accomplished,
and the process is returned to the previous step 901.

In a step 904, a loop to which a node belongs
15 is investigated, and a "data" designation is
investigated based upon the loop table 210. In a step
905, a check is made as to whether or not the "data"
designation corresponds to "L1". When the "data"
designation corresponds to "L1", the process is
20 advanced to a step 909. In this step 909, a value of a
latency of the L1 cache is set to an edge. Then, the
process is returned to the previous step 901. When the
"data" designation is not equal to "L1", the process is
advanced to a step 906. In this step 906, a check is
25 made as to whether or not the "data" designation
corresponds to "L2". When the "data" designation
corresponds to "L2", the process is advanced to a step
908. In this step 908, a value of a latency of the L2

cache is set to an edge, and then, the process is returned to the previous step 901. When the "data" designation is not equal to "L2", namely, in the case that the "data" designation corresponds to the "MEM",
5 or no designation, the process is advanced to a step 907. In this step 907, a value of a latency of the main storage access operation is set to the edge, and then, the process is returned to the previous step 901.

Fig. 8 represents an example of such a case
10 where a "data" designation corresponding to a node of each load is "L2", and a latency of the L2 cache is 10 cycles.

This embodiment can achieve the following effect. That is, since a latency of a DAG edge is set
15 in response to such a fact that subject data of a load instruction is present in any one of the L1 cache, the L2 cache, and the main storage, the instruction scheduling method can be applied in which the more correct latency is assumed with respect to the load
20 instruction.

Next, description is made of such an example that a prefetch optimizing operation is carried out in the memory optimizing process 502 according to the present invention. Fig. 10 is a flowchart for
25 explaining a process sequence of the prefetch optimizing method.

In a step 1001 of the prefetch optimizing processes shown in Fig. 10, judgement is made as to

whether or not there is a loop which has not yet been processed. When there is an unprocessed loop, the unprocessed loop is derived in a step 1002. In a step 1003, a "data" designation of the derived loop is
5 investigated based upon the loop table 210. When the "data" designation corresponds to "L1", the process is returned to the previous step 1001. This is because it is so judged that there is no effect in the prefetch optimizing method, since the data of this loop is
10 already present in the L1 cache. In a step 1004, a check is made as to whether or not the "data" designation corresponds to "L2". When the "data" designation is equal to "L2", the process is advanced to a step 1006. In this step 1006, the latency of the
15 L2 cache is set to "L". When the "data" designation is not equal to "L2", namely in a case where the "data" designation corresponds to either "MEM" or no designation, the process is advanced to a step 1005. In this step 1005, the latency of the main storage is
20 set to "L".

In a step 1007, the total number of execution cycles per loop iteration is calculated, and then, the calculated execution cycle number is set to "C". In a step 1008, such a value obtained by rounding up (L/C)
25 to an integer is set to "N". In a step 1009, a prefetch code after N iterations is produced, and this produced prefetch code is inserted into the loop.

This embodiment has such an effect that since

the latency of the main storage is set in response to such a fact as to whether or not the data is present in the L2 cache, the prefetch code is produced based upon the more proper prefetch distance. Also, this
5 embodiment owns such an effect that when the data is present in the L1 cache, the unnecessary prefetch code is not produced.

Next, description is made of such an example that the tiling method is carried out in the memory
10 optimizing process 502 according to the present invention. Fig. 11 shows a flowchart for explaining a sequential process as to the tiling method.

In a step 1101 of the tiling method shown in Fig. 11, judgement is made as to whether or not there
15 is a loop which has not yet been processed. When there is an unprocessed loop, the unprocessed loop is derived in a step 1103. In a step 1103, a "data" designation of the derived loop is investigated based upon the loop table 210. When the "data" designation corresponds to
20 "L1", the process is returned to the previous step 1101. This is because it is so judged that there is no effect in the tiling method, since the data of this loop is already present in the L1 cache. In a step 1104, a check is made as to whether or not the "data"
25 designation corresponds to "L2". When the "data" designation is equal to "L2", the process is advanced to a step 1106. In this step 1106, a target cache is set as "L1". When the "data" designation is not equal

to "L2", namely in a case where the "data" designation corresponds to either "MEM" or no designation, the process is advanced to a step 1105. In this step 1105, a target cache is set as "L2".

5 In a step 1107, an application condition of the tiling method is investigated, and it is determined as to whether or not an unprocessed loop corresponds to a target loop to which the tiling method is applied. As the application condition, a check is made as to
10 whether or not the unprocessed loop corresponds to a multiply nested loop, whether or not the unprocessed loop can satisfy a dependence test, or whether or not the unprocessed loop can achieve the tiling effect. A detailed judging method of the application condition is
15 described in Publication 2. When the unprocessed loop is not equal to the target loop to which the tiling method is applied, the process of this loop is accomplished. Then, the process is returned to the previous step 1101. To the contrary, when the
20 unprocessed loop corresponds to the target loop to which the tiling method is applied, the process is advanced to a step 1108.

 In the step 1108, a tile size is determined based upon the cache size, an associative degree, and
25 the like of the target cache determined in both the steps 1105 and 1106. A method of determining a tile size is described in Publication 2.

 In a step 1109, a tiling conversion process

is carried out in accordance with the tile size determined in the step 1108. The tiling conversion process is also described in Publication 2.

In accordance with this embodiment, since the
5 target cache is determined according to the "data" designation, there is such an effect that the tiling conversion based upon the more proper tile size is applied. Also, there is another effect that in a case where the data is present in the L1 cache, unnecessary
10 tiling conversion is not applied.

Next, description is made of such an example that both a loop interchange and a loop unrolling are carried out in the memory optimizing process 502 according to the present invention. Fig. 12 shows a
15 flowchart for explaining sequential processes as to both the loop interchange and the loop unrolling.

In accordance with this embodiment, in a case where data is present in either the L1 cache or the L2 cache, it is so considered that there is no effect
20 achieved by the loop interchange and the loop unrolling. Thus, the application of the loop interchange and the loop unrolling is ceased. Only when it is predictable that data is present in the main storage, the memory optimizing process 502 is applied.

25 In a step 1201 of the loop interchange and the loop unrolling shown in Fig. 12, judgement is made as to whether or not there is a loop which has not yet been processed. When there is an unprocessed loop, the

unprocessed loop is derived in a step 1202. In a step 1203, a "data" designation of the derived loop is investigated based upon the loop table 210. When the "data" designation corresponds to "L1", or "L2", the process is returned to the previous step 1201. This is because it is so judged that there is no effect in the optimizing process since the data of this loop is already present in the cache. In a step 1204, an application condition of the loop interchange is investigated. When the unprocessed loop corresponds to the application target loop, the loop interchange is applied to this unprocessed loop in a step 1205. In a step 1206, the application condition of the loop unrolling is investigated. When the unprocessed loop corresponds to the application target loop, the loop unrolling is applied to this unprocessed loop in a step 1207.

This embodiment owns the following effect. That is, in such a case that the data contained in the loop is mainly present in the cache, both the loop interchange and the loop unrolling are not applied when they are unnecessary.

Alternatively, all or several optimizing processes of the memory optimizing processes 502 may be combined with each other. Furthermore, any one optimizing process contained in the memory optimizing processes 502 may be carried out.

It should be further understood by those

skilled in the art that although the foregoing
description has been made on embodiments of the
invention, the invention is not limited thereto and
various changes and modifications may be made without
5 departing from the spirit of the invention and the
scope of the appended claims.